

Network Addressing Notes

What's a node number?

Apr 2, 1990

These notes explore the conventions used in network addressing for the VME Local Stations. There are two sets of network protocols that are handled by the local stations. The "classic" protocols do not use the Acnet header, while both the DZero and Acnet protocols do.

Acnet header protocols—current

When a request is received, the node# byte is taken from the source node/lan word and used to install a copy of the 6-byte network address in the Node Address Table (using the NATENTER routine). This node# is used internally for the destination node for the reply. This method means that any network address can be used for any node, just as long as the requests from different nodes (with different network addresses) do not use the same node#. If they did, the second one's network address would overwrite that of the first one in the NAT. That could mean that the replies from an active request from the first node would all-of-a-sudden be sent to the second node instead. Note that the lan byte is not checked here; however, it is returned in the reply as it was received in the request.

If the local station wants to send a request or USM, it starts with the node# as a destination. The corresponding entry is checked in the NAT. If it is there, then the stored network address is used. If it is not there, then the local station's address is used, with the node# replacing the sixth byte.

Classic protocol—current

There is no source node/lan information in the message. Instead, the sixth byte of the source network address is interpreted as a node#. For any message, the NAT entry is updated. Therefore, the network address used by a node for this protocol is not arbitrary. The last byte must be unique. The other 5 bytes can be used arbitrarily for requesters. But when a local station wants to talk to a node given a node#, it will still assume the network address is remarkably similar to the local station address in the case that the NAT entry is empty.

Acnet console practice

The consoles run by accelerator operators are all downloaded with a common logical node table that contains all the network addresses that are assigned to each logical node. The logical node numbers are 16-bit values found in the database. The values are in the range 0x0001–0x03FF, although only the first 72 values are currently used. In some cases, more than one logical node can refer to the same physical node. When a request is received from one of these consoles, the NAT table should be used to keep the network address. Seventy-two entries would be sufficient for the present.

For front end nodes, the current practice is to hand out blocks of 256 node numbers for specific uses. For these cases, the lan byte of the source lan-node word is greater than three. Thus, the new QPM systems have been assigned the "lan" byte value of 4. The new Linac stations may be assigned the value of 5, for example. DZero stations might be assigned the value of 6, if desired. To simplify the handling of micro-based stations,

the lan-node word is used as the last two bytes of the network address, the first four bytes being constant. The value of the constant is 0x40020000. Therefore, when a request is received, the network address does not have to be saved in a table. Only the source lan-node is saved as the destination lan-node for the reply. When the frame is to be transmitted, the value of the destination network address is simply the constant value for the first 4 bytes and the destination lan-node used for the last 2 bytes.

Local station lan byte

When a local station issues a request of a USM Acnet header-based message, the source lan-node must be filled in for the Acnet header. One way to do this is to use a constant value known to the program. Another way is to keep it in non-volatile memory in the TRING table. Either way, it would be the same value for each request or USM message. (This does not apply for the reply message, since whatever was received in the request is simply echoed in the reply.) This method will mean that the user does not have to enter the lan byte value when she types in a channel number, for example, since it is fixed and supplied automatically.

Lan byte in ident

When a request is received that gives an ident for which data is requested, the lan byte will be a value like 5, for example, if it is a Linac device. The code that checks for a match on the node will have to work for this case. Perhaps the MYNODE word will have to be a value like 0x0508, for example. The hi byte would be the local station's lan number. Or, the code could perhaps ignore the given lan byte value and only check on the node byte. Right now, the first word of long idents is matched against the whole word stored at MYNODE (A5).